H4 GRANT
IN-62-CR
174549
17P.

Fault Diagnosis in Sparse Multiprocessor Systems*

Douglas M. Blough, Gregory F. Sullivan, and Gerald M. Masson
Department of Computer Science
The Johns Hopkins University
Baltimore, MD 21218

Abstract

In this paper, the problem of fault diagnosis in multiprocessor systems is considered under a uniformly probabilistic model in which processors are faulty with probability p. This work focuses on minimizing the number of tests that must be conducted in order to correctly diagnose the state of every processor in the system with high probability. A diagnosis algorithm that can correctly diagnose the state of every processor with probability approaching one in a class of systems performing slightly greater than a linear number of tests is presented. A nearly matching lower bound on the number of tests required to achieve correct diagnosis in arbitrary systems is also proven. The number of tests required under this probabilistic model is shown to be significantly less than under a bounded-size fault set model. Because the number of tests that must be conducted is a measure of the diagnosis overhead, these results represent a dramatic improvement in the performance of system-level diagnosis techniques.

1 Introduction

In this paper, the fault diagnosis capabilities of multiprocessor systems in the presence of permanently faulty processors are examined. This problem has been well studied under the assumption that the number of faulty processors in the system is bounded by some value t. It has been shown that nt tests are necessary and sufficient to correctly diagnose a system of n processors in this situation [1]. The results of this paper will show that under a probabilistic model in which processors are faulty with probability p independently of one another that correct diagnosis

* This work was supported by NASA under Grant NSG-1442.

(NASA-CR-183396) FAULT DIAGNOSIS IN SPARSE MULTIPROCESSOR SYSTEMS Semiannual Progress Report (Johns Hopkins Univ.) 17 p CSCL 09B

N89-12282

Unclas
G3/62 0174549

$O(n \cdot \omega(n))$ tests where $\omega(n) \to \infty$ (arbitrarily slowly) as $n \to \infty$. Thus, in the bounded-size fault set model a quadratic number of tests are required to diagnose a linear number of faults while under this probabilistic model a linear expected number of faults can be diagnosed with high probability using a number of tests growing slightly faster than n.

The problem of multiprocessor system diagnosis in the presence of permanent faults has been addressed from a probabilistic viewpoint in several papers [2,3,4]. The first paper concerning probabilistic diagnosis [2] examined heterogeneous systems in which each processor has an associated probability of failure. The authors examined the class of systems known as p-probabilistically diagnosable systems in which any fault set that has probability greater than or equal to p of occurring is uniquely diagnosable. The problem of determining whether a given system is p-probabilistically diagnosable has been shown to be co-NP-complete [3] while an $O(n^3)$ algorithm has been given [4] for determining the most likely fault set of a system in the closely related weighted model.

In p-probabilistically diagnosable systems, fault sets with probability of occurrence slightly less than p can exist. Hence, the most likely fault set may be only slightly more probable than the next most likely fault set, meaning that the probability of choosing the wrong fault set may be relatively high. In [5], the author examined systems for which the correct fault set can be identified with high probability. The model utilized applies to homogeneous systems in which each processor has a common probability of failure p. An efficient diagnosis algorithm was presented that correctly diagnoses a class of systems containing $cn \log n$ tests, for $c > 1/(\log 1/p)$, with probability approaching one.

It was also claimed in [5] that this result was the best possible, i.e. that all algorithms must have probability approaching zero of achieving correct diagnosis in systems containing $o(n \log n)$ tests. Unfortunately, due to a subtle flaw in the proof, this result is untrue. This result was also used in [6] to prove a similarly flawed lower bound in a more general probabilistic model. A counterexample to the lower bound in [5] is given in which correct diagnosis is achieved with constant probability in a sequence of digraphs containing $n - 1$ tests. Also in this paper a diagnosis algorithm that produces correct diagnosis with probability approaching one in digraphs containing slightly more than a linear number of tests is given. Finally, a nearly matching lower bound on the number of tests required to achieve correct diagnosis with probability approaching one is proven.

2 Preliminaries

The fundamental multiprocessor system model utilized in this paper was proposed in [7]. In this model a system is represented as a directed graph with vertices of the digraph representing processors in the system and edges of the digraph representing tests performed by one processor on another processor. In this section, all basic quantities related to this model are defined and methods of diagnosis algorithm performance evaluation are examined.

2.1 Basic Definitions

For a system composed of n processors, the set of processors will be represented by $U = \{u_1, \ldots, u_n\}$. It is assumed that these processors are capable of performing tests on one another. This situation will be represented by a digraph $G(U, E)$, where the vertex set U corresponds to the set of processors of the system and $(u, v) \in E$ if and only if processor u tests processor v in the system. Associated with each $(u, v) \in E$ is a *test outcome*. This outcome will be a 1(0) if u evaluates v as faulty (fault-free). A complete collection of test outcomes constitutes a *syndrome*. Below syndromes, fault sets, and other fundamental concepts are formally defined.

Definition 1 *For a digraph $G(U, E)$, a* syndrome *is a function from E to $\{0, 1\}$.*

Definition 2 *For a digraph $G(U, E)$, a* fault set *is a subset of the vertex set U.*

Definition 3 *For a digraph $G(U, E)$ and $u \in U$, the* tester set *of u, denoted by $\Gamma^{-1}(u)$, is given by*

$$\Gamma^{-1}(u) = \{v \in U : (v, u) \in E\}$$

Definition 4 *For a digraph $G(U, E)$, a syndrome S, and $u \in U$, the* failure set *of u, denoted by $\Delta_{in}(u)$, is given by*

$$\Delta_{in}(u) = \{v \in \Gamma^{-1}(u) \quad S((v, u)) = 1\}$$

2.2 Diagnosis Algorithm Evaluation

A fundamental problem in multiprocessor systems is to identify the faulty processors in a system given a syndrome. An algorithm for this problem is referred to as a *diagnosis algorithm*. In much of the previous work in the system-level diagnosis area, diagnosis algorithm evaluation has focused on worst-case performance. Under a bounded-size fault set model, correct diagnosis can be guaranteed if the number of faulty processors in the system is no greater than some value $t < n/2$. Since this bound can only be satisfied with a given probability, a better measure of diagnosis algorithm performance is the probability that it correctly identifies the faulty processors in the system under a probabilistic model for the faults and test outcomes in a system. Such a model is presented in this paper.

A diagnosis algorithm takes a syndrome as input and outputs a subset of the processors in the system. This subset contains exactly the processors diagnosed as faulty by the algorithm. Thus, for a set of faulty processors and a syndrome it is possible to evaluate if the output of a deterministic algorithm is correct by comparing the algorithm's output with the set of faulty processors. Syndrome, fault set pairs are therefore used as the basic element in the subsequent probabilistic analysis of diagnosis algorithm performance. Before proceeding with this analysis however, the notion of correct diagnosis must be formally defined. For a syndrome S from a digraph $G(U, E)$, and a deterministic algorithm A, let

$$Faulty_A(S) = \{u \in U : \text{Algorithm } A \text{ diagnoses } u \text{ as faulty when run on } S\}$$

Thus, $Faulty_A(S)$ represents the output of Algorithm A when run on syndrome S. With this, the diagnosis of an algorithm on a syndrome, fault set pair is characterized in Definition 5.

Definition 5 *For a syndrome, fault set pair (S, F) from a digraph $G(U, E)$, a deterministic algorithm A is said to produce*

correct diagnosis *if and only if* $Faulty_A(S) = F$,

partial diagnosis *if and only if* $Faulty_A(S) \subset F$, *and*

false alarm diagnosis *if and only if* $Faulty_A(S) \not\subseteq F$.

Note that Definition 5 differs from that used in some previous work where correct diagnosis may include faulty processors that are identified as fault-free so long as no fault-free processor is identified as faulty. In Definition 5, diagnosis is correct only when each fault-free processor is identified as fault-free *and* each faulty processor is identified as faulty.

3 Probabilistic Model

In this section, a probabilistic model for the behavior of a multiprocessor system is presented. In this model, processors are faulty with probability p, fault-free processors always produce the correct outcome when performing a test, and no assumptions are made concerning the outcomes of tests performed by faulty processors. It will be shown in this paper that in contrast to the bounded-size fault set model, correct diagnosis can be achieved with high probability in this model at relatively low cost.

For a digraph $G(U, E)$, the sample space of this probability model will consist of all syndrome, fault set pairs in that digraph. Formally,

$$\Omega_{G(U,E)} = \{(S, F) : F \subseteq U \text{ and } S \text{ is a function from } E \text{ to } \{0, 1\}\}.$$

Since no assumptions have been made concerning the outcomes of tests performed by faulty processors, the probability of a particular syndrome given a fault set may not be specified in this model. The basic events of the model consist of sets of syndrome, fault set pairs which have the same fault set and whose syndromes are identical except for the labels on edges out of faulty processors. Formally, a syndrome, fault set pair (S', F') is contained in a basic event B defined as follows:

$$B = \{(S, F) : F = F' \text{ and } \forall(u, v) \in E \text{ with } u \in U - F,\ S((u, v)) = S'((u, v))\}$$

Note that there is a unique fault set associated with each basic event but that each event may contain many distinct syndrome, fault set pairs. Now, let

$$\mathcal{B}_{G(U,E)} = \{B : B \text{ is a basic event of } G(U, E)\}.$$

The family of events $\mathcal{F}_{G(U,E)}$ in this probability space is the set of all subsets of $\mathcal{B}_{G(U,E)}$. For a basic event B from a digraph $G(U, E)$, let

$$\begin{aligned} E_{c0} &= \{(u, v) \in E : \forall(S, F) \in B,\ u \in U - F \text{ and } v \in U - F\} \text{ and} \\ E_{c1} &= \{(u, v) \in E : \forall(S, F) \in B,\ u \in U - F \text{ and } v \in F\}. \end{aligned}$$

These sets represent respectively, the set of edges that must be labeled zero (fault-free processors testing fault-free processors) and the set of edges that must be labeled one (fault-free processors testing faulty processors). Given these sets, the probability of a basic event B in a digraph $G(U, E)$ is defined as follows:

$$P_G(B) = \begin{cases} 0 \text{ if } \exists(u, v) \in E_{c0} \text{ s.t. } \forall(S, F) \in B,\ S((u, v)) = 1 \quad \text{or} \\ \qquad \exists(u, v) \in E_{c1} \text{ s.t. } \forall(S, F) \in B,\ S((u, v)) = 0 \\ p^{|F'|}(1 - p)^{n-|F'|} \text{ otherwise} \end{cases}$$

where F' represents the unique fault set associated with B. The condition for which a basic event has zero probability of occurrence is simply a check to make sure that no fault-free processor produces an incorrect test outcome. Clearly,

$$\sum_{B \in \mathcal{B}} P_G(B) = 1$$

and, hence, this is a legitimate probability measure.

The primary measure of the performance of a diagnosis algorithm used in this paper will be the probability that the algorithm produces correct diagnosis as defined in Definition 5. For a digraph $G(U, E)$ and a deterministic algorithm A, let

$$\text{Correct}_G(A) = \{(S, F) : Faulty_A(S) = F\}$$

and let $\text{NotCorrect}_G(A)$ represent the complement of $\text{Correct}_G(A)$. Thus, $\text{Correct}_G(A)$ represents the set of all syndrome, fault set pairs in a digraph for which Algorithm A produces correct diagnosis. Note that it may be the case that $\text{Correct}_G(A) \notin \mathcal{F}_G$ in which case $P_G(\text{Correct}_G(A))$ will not be defined. The output of a particular diagnosis algorithm may depend on the outcomes of tests performed by faulty processors and thus, the probability of correct diagnosis for the algorithm cannot be determined until a probability distribution on these edges is specified.

For a digraph $G(U, E)$, let P_G' be a probability function defined on Ω_G such that the family of events is equal to all subsets of Ω_G and $\forall B \in \mathcal{B}_G$, $P_G'(B) = P_G(B)$. Such a probability function will be referred to as a *refinement* of P_G. Now, let $\mathcal{P}_G$ represent the set of all refinements of P_G. Since any type of behavior of the faulty processors is allowed in this model, the probability of correct diagnosis for a deterministic algorithm A in a digraph $G(U, E)$, denoted by $\text{PCD}_G(A)$ is defined to be

$$\text{PCD}_G(A) = \min_{P_G' \in \mathcal{P}_G} P_G'(\text{Correct}_G(A)) = \min_{P_G' \in \mathcal{P}_G} \sum_{(S,F) \in \text{Correct}_G(A)} P_G'((S, F))$$

Thus, when calculating the probability of correct diagnosis for an algorithm it is assumed that the faulty processors perform their tests in the manner most detrimental to the algorithm. Given a syndrome S, a random diagnosis algorithm A chooses a fault set F with some probability call it $p_{A,S}(F)$ where $\sum_{F \subseteq U} p_{A,S}(F) = 1$. Thus, for a digraph $G(U, E)$ and a random diagnosis algorithm A, the probability of correct diagnosis for Algorithm A is defined to be

$$\text{PCD}_G(A) = \min_{P_G' \in \mathcal{P}_G} \sum_{(S,F) \in \Omega_G} P_G'((S, F)) \cdot p_{A,S}(F)$$

4 Diagnosis Below n log n Edges

In [5], a powerful and efficient diagnosis algorithm that achieves correct diagnosis with probability approaching one in sequences of digraphs containing $cn \log n$ edges, for $c > 1/(\log 1/p)$, was presented. In this section, the question of whether correct diagnosis is possible in digraphs containing $o(n \log n)$ edges is considered. In particular, a sequence of digraphs containing $n-1$ edges is exhibited for which a simple diagnosis algorithm can achieve correct diagnosis with constant probability.

Consider a sequence of digraphs $G_n(U_n, E_n)$ with $U_n = \{u_1, \ldots, u_n\}$ and E_n defined as follows:

$$E_n = \{(u_1, u_2), (u_1, u_3), \ldots, (u_1, u_{n-1}), (u_1, u_n)\},$$

i.e. u_1 tests all other processors. Now, consider the following simple diagnosis algorithm.

Algorithm Naive

Input: A syndrome S in a digraph $G(U, E)$.
Output: A set $F \subseteq U$.

$F \leftarrow \emptyset$
for each $v \in \{u_2, u_3, \ldots, u_n\}$
    if $S((u_1, v)) = 1$ then $F \leftarrow F \cup \{v\}$.

Algorithm Naive simply assumes that u_1 is fault-free and diagnoses a processor as faulty if and only if it is failed by u_1. Clearly, if u_1 is faulty, Algorithm Naive incorrectly diagnoses u_1 itself. If u_1 is fault-free however, Algorithm Naive produces correct diagnosis. Thus, $\forall {P_{G_n}}' \in \mathcal{P}_{G_n}$

$$\begin{aligned} {P_{G_n}}'(\text{Correct}_{G_n}(\text{Naive})) &= {P_{G_n}}'(\{(S, F) : u_1 \text{ is fault-free}\}) \\ &= 1 - p \end{aligned}$$

and therefore

$$\text{PCD}_{G_n}(\text{Naive}) = 1 - p.$$

Thus, this simple diagnosis algorithm produces correct diagnosis with constant probability in a sequence of digraphs containing exactly $n-1$ edges.

The digraphs of the given sequence are composed of one processor testing the remaining processors. It will be shown that this highly irregular structure whereby some processors conduct a large number of tests while others may not conduct

any is common to all systems of $o(n \log n)$ edges that can be diagnosed with high probability. In Section 6, a class of irregular digraphs possessing a number of edges growing just faster than n is given for which correct diagnosis can be achieved with probability approaching one. In Section 7, it is shown that a linear number of edges is required to achieve correct diagnosis with high probability in arbitrary digraphs.

5 A Simple Majority–Vote Algorithm

In this section, a simple yet powerful diagnosis algorithm known as Algorithm Majority is presented. In Algorithm Majority a processor is diagnosed as faulty if and only if it is failed by more than 1/2 the processors in its tester set.

Algorithm Majority

Input: A syndrome S in a digraph $G(U, E)$.
Output: A set $F \subseteq U$.

$F \leftarrow \emptyset$
for each $u \in U$
  if $|\Delta_{in}(u)| > \frac{|\Gamma^{-1}(u)|}{2}$ then $F \leftarrow F \cup \{u\}$

Theorem 1 *For a digraph $G(U, E)$, Algorithm Majority has a time complexity of $O(|E|)$ and a space complexity of $O(|E|)$.*

Proof: The failure set cardinalities as well as the tester set cardinalities can be calculated in a single traversal of the labeled adjacency lists of the digraph. This requires $O(|E|)$ time. The only storage requirement for the algorithm aside from the input and output is temporary variables to hold these values as they are calculated. Hence, the space complexity is also $(|E|)$. ∎

Algorithm Majority is slightly more sophisticated than Algorithm Naive. Rather than blindly believing the test outcomes of a single processor, it relies on a majority-vote among the processors in the tester set of a given processor. It should be noted that for the special class of systems in which one processor tests every other processor and no other tests are conducted, Algorithms Naive and Majority are equivalent. Intuitively, when $p < 1/2$ the majority of processors in the system are fault-free and Algorithm Majority should correctly diagnose most of the processors in the system. In the next section, the performance of Algorithm Majority is considered in detail.

6 Performance of Algorithm Majority

In this section, it is shown that for a class of irregularly structured systems utilizing a number of tests growing just faster than n, Algorithm Naive correctly diagnoses every processor with probability approaching one. The exact number of tests required by Algorithm Majority to achieve a given probability of correct diagnosis on systems in this class is also examined.

6.1 Asymptotic Results

Consider a class of systems in which there is a set of processors known as the *testers*. The systems are such that any processor which is a tester tests all other processors in the system. Any processor that is not a tester conducts no tests. Thus, a (small) fraction of the processors are relied upon to satisfy all the testing requirements of the system. Such a digraph will be referred to as a *tester digraph*, formally defined below.

Definition 6 *A digraph $G(U, E)$ is said to be a* tester digraph *if and only if $\exists T_G \subseteq U$ such that*

$$E = \{(u, v) : u \in T_G, v \in U, \text{ and } u \neq v\}.$$

The set T_G is known as the testing set *of G.*

Figure 1 is an example of a tester digraph with 3 testers and 8 vertices. Assume that more than 1/2 the testers in a tester digraph are fault-free. Clearly, more than 1/2 the tests conducted on any processor that is not a tester will be accurate and each such processor will be correctly diagnosed by Algorithm Majority. Now, consider any tester t. If t is faulty, more than 1/2 the processors testing it are fault-free and will fail it, meaning that t will be correctly diagnosed by Algorithm Majority. If t is fault-free, at least 1/2 the processors testing it are fault-free and will pass it. Since t is not failed by a majority of its tester set, it will again be correctly diagnosed by Algorithm Majority. Hence, if more than 1/2 the testers in a tester digraph are fault-free, Algorithm Majority produces correct diagnosis. Theorem 2 shows that if the number of testers is given by any function that increases with n, this condition will be achieved with probability approaching one and hence the probability of correct diagnosis for Algorithm Majority approaches one. In order to prove this result the following corollary [8] to a theorem proved by Chernoff [9] is needed.

Corollary 1 *Let Y be a binomial random variable with parameters n and p. Then*

$$P(Y \leq cnp) \leq e^{-(1-c)^2 np/2}, \quad 0 < c \leq 1$$

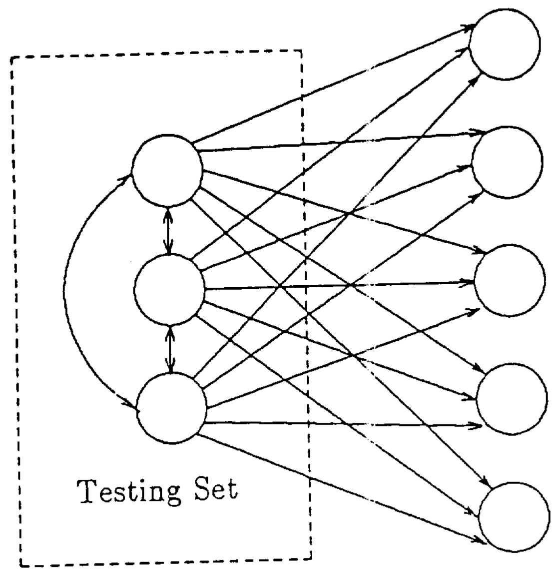

Figure 1: A Tester Digraph

$$P(Y \geq cnp) \leq e^{-(c-1)^2 np/3}, \quad c \geq 1$$

Theorem 2 *Let $G_n(U_n, E_n)$ be a sequence of tester digraphs on n vertices with testing sets T_{G_n} satisfying $|T_{G_n}| = \omega(n)$, where $\omega(n) \to \infty$ as $n \to \infty$. If $p < 1/2$, then* $\mathrm{PCD}_{G_n}(\mathrm{Majority}) \to 1$ *as $n \to \infty$.*

Proof: Let

$$\begin{aligned} \mathrm{GoodMaj}_{G_n} &= \{(S,F) : |T_{G_n} \cap (U_n - F)| > \frac{|T_{G_n}|}{2} \text{ and } \forall (u,v) \in E_{c0}, \\ & \qquad S((u,v)) = 0 \text{ and } \forall (u,v) \in E_{c1}, S((u,v)) = 1\} \end{aligned}$$

Clearly,

$$\mathrm{GoodMaj}_{G_n} \subseteq \mathrm{Correct}_{G_n}(\mathrm{Majority})$$

and therefore, $\forall {P_{G_n}}' \in \mathcal{P}_{G_n}$

$$\begin{aligned} {P_{G_n}}'(\mathrm{Correct}_{G_n}(\mathrm{Majority})) &\geq {P_{G_n}}'(\mathrm{GoodMaj}_{G_n}) \\ &= 1 - \sum_{i=0}^{\left\lfloor \frac{|T_{G_n}|}{2} \right\rfloor} \binom{|T_{G_n}|}{i} (1-p)^i p^{|T_{G_n}|-i} \end{aligned}$$

Now, since $p < 1/2$
$$\frac{|T_{G_n}|}{2} = c(1-p)|T_{G_n}|,\ c < 1$$

and thus by Corollary 1,

$$\begin{aligned} P_{G_n}{}'(\text{Correct}_{G_n}(\text{Majority})) &\geq 1 - \left[e^{-(1-c)^2/2}\right]^{(1-p)\omega(n)} \\ &\rightarrow 1 \end{aligned}$$

Therefore
$$\text{PCD}_{G_n}(\text{Majority}) \rightarrow 1.$$

∎

Thus, Algorithm Majority produces correct diagnosis with probability approaching one in a class of digraphs containing a number of edges given by $n \cdot \omega(n)$, where $\omega(n)$ is any function that increases with n. This is an extremely promising result because under a bounded-size fault set model a quadratic number of tests are required to withstand a linear number of faults while this shows that in this probabilistic model a linear expected number of faults can be tolerated with a number of tests that is arbitrarily close to linear.

6.2 Concrete Bounds

In this section, the number of tests required to achieve a given probability of correct diagnosis in tester digraphs using Algorithm Majority is examined. For a tester digraph $G(U, E)$ with testing set T_G

$$\text{PCD}_G(\text{Majority}) \geq \sum_{i=0}^{\left\lfloor \frac{|T_G|}{2} \right\rfloor} \binom{|T_G|}{i} (1-p)^i p^{|T_G|-i} \qquad (1)$$

Note that the probability of correct diagnosis depends only on the testing set cardinality and not on n. For a given probability of failure, Inequality 1 can be used to determine the number of testers needed for Algorithm Majority to achieve a specific probability of correct diagnosis. The size of the testing set required to achieve a correct diagnosis probability of 0.99 for various values of p is shown in Table 1. If the probability of failure of a processor is 0.01, Algorithm Majority can achieve correct diagnosis with a probability of 0.99 using a single test per processor regardless of the number of processors in the system. This corresponds exactly to the example given in Section 4 where a single tester tests every other processor in the system. Hence, the total number of tests utilized in this situation is $n - 1$. For a probability of failure of 0.1 the tester set need only be of cardinality 5 for Algorithm Majority

p	$\lvert T_G \rvert$
0.01	1
0.05	3
0.10	5
0.20	13
0.30	31
0.40	133

Table 1: Size of Testing Set Required for Correct Diagnosis Probability of 0.99

n	p	Bounded-size	Probabilistic
100	0.01	400	99
100	0.10	1800	495
100	0.30	4100	3069
1000	0.01	18000	999
1000	0.10	123000	4995
1000	0.30	334000	30969

Table 2: Total Number of Tests Necessary for Correct Diagnosis Probability of 0.99

to achieve a probability of correct diagnosis of 0.99. Thus, when the probability of failure is small correct diagnosis can be achieved with high probability using a total number of tests that is near n. When p is near 1/2, more tests are necessary. Since nearly 1/2 the processors in the system will be faulty in this situation it is to be expected that a larger number of tests are required. The important point is that the total number of tests remains proportional to n regardless of the value of p.

These results can be compared with the number of tests required under the bounded-size fault set model in the following manner. For a given n and p, determine t such that the probability of more than t out of the n processors being faulty is no greater than 0.01. Table 2 shows the results of this comparison for various values of n and p. For large n and small p the number of tests required under the probabilistic

model is dramatically lower than the number required under the bounded-size fault set model. For example, when $n = 1000$ and $p = 0.01$, the number of tests required in the probabilistic model is reduced by a factor of 18 over the bounded-size fault set model.

7 A Lower Bound on the Number of Tests Necessary for Correct Diagnosis

In this section, a lower bound on the number of tests necessary to achieve correct diagnosis with high probability is proven. It is shown that if the number of edges in an arbitrary sequence of digraphs grows slower than n then all diagnosis algorithms have probability approaching zero of achieving correct diagnosis. This result implies that Algorithm Majority achieves a probability approaching one of correct diagnosis on systems that are very nearly as sparse as possible. Thus, this relatively simple diagnosis algorithm is indeed extremely powerful.

When the number of edges in a sequence of digraphs grows slower than n, isolated processors must exist. Intuitively, no diagnosis algorithm should be capable of correctly identifying the state of all these isolated processors with high probability, making diagnosis in such situations impossible. This is formally proven in the Theorem 3. The essence of the proof of Theorem 3 can be explained quite simply. To prove that a deterministic diagnosis algorithm A has a probability approaching zero of achieving correct diagnosis in a sequence of digraphs $G_n(U_n, E_n)$, a set of (S, F) pairs disjoint from $\mathrm{Correct}_{G_n}(A)$ must be exhibited that has a probability dominating the probability of $\mathrm{Correct}_{G_n}(A)$. For a given syndrome with isolated vertices, it can be shown that so long as the number of isolated vertices approaches infinity, the probability of that syndrome and a fault set with a particular labeling of the isolated vertices is dominated by the probability of that syndrome and the fault sets in which the isolated processors are relabeled in all possible ways. Thus, for any $(S, F) \in \mathrm{Correct}_{G_n}(A)$, a set of syndrome, fault set pairs disjoint from $\mathrm{Correct}_{G_n}(A)$ can be exhibited that has probability dominating the probability of (S, F). It is also shown that there exists a deterministic diagnosis algorithm that has performance at least as good as the performance of any random algorithm, thus completing the proof.

Theorem 3 *Let $G_n(U_n, E_n)$ be a sequence of digraphs on n vertices with $0 < p < 1$ and $|E_n| \in o(n)$. For any random or deterministic diagnosis algorithm A, $\mathrm{PCD}_{G_n}(A) \to 0$ as $n \to \infty$.*

Proof: Assume $\exists n_0, c > 0$, and a deterministic algorithm A such that $\forall n \geq n_0$,

$\mathrm{PCD}_{G_n}(A) \geq c$. This implies that $\forall P_{G_n}{}' \in \mathcal{P}_{G_n}$ and $\forall n \geq n_0$,

$$P_{G_n}{}'(\mathrm{Correct}_{G_n}(A)) \geq c.$$

Now, let $ISO_{G_n} \subseteq U_n$ represent the set of isolated vertices in $G_n(U_n, E_n)$. Clearly,

$$|ISO_{G_n}| \geq n - 2|E_n| \to \infty.$$

For a syndrome, fault set pair $(S, F) \in \mathrm{Correct}_{G_n}(A)$ let

$$\mathrm{Relabel}_{(S,F)} = \{(S', F') : S' = S, F' \neq F, \text{ and } F - ISO_{G_n} = F' - ISO_{G_n}\}$$

and let

$$\mathrm{AllLabel}_{(S,F)} = \mathrm{Relabel}_{(S,F)} \cup \{(S, F)\}.$$

Thus, $\mathrm{Relabel}_{(S,F)}$ consists of the syndrome, fault set pairs in which the processors of ISO_{G_n} are relabeled in all possible ways. Clearly, $\forall P_{G_n}{}' \in \mathcal{P}_{G_n}$

$$\begin{aligned}
&P_{G_n}{}'(\mathrm{NotCorrect}_{G_n}(A)) \\
&\geq \sum_{(S,F)\in \mathrm{Correct}_{G_n}(A)} P_{G_n}{}'(\mathrm{Relabel}_{(S,F)}) \\
&= \sum_{(S,F)\in \mathrm{Correct}_{G_n}(A)} \left[P_{G_n}{}'(\mathrm{AllLabel}_{(S,F)}) - P_{G_n}{}'((S,F))\right]
\end{aligned}$$

and since all processors in the set ISO_{G_n} are isolated,

$$P_{G_n}{}'((S,F)) = p^{|ISO_{G_n}\cap F|}(1-p)^{|ISO_{G_n}\cap(U_n-F)|}P_{G_n}{}'(\mathrm{AllLabel}_{(S,F)}).$$

Therefore, $\forall P_{G_n}{}' \in \mathcal{P}_{G_n}$

$$\begin{aligned}
&\sum_{(S,F)\in \mathrm{Correct}_{G_n}(A)} P_{G_n}{}'(\mathrm{AllLabel}_{(S,F)}) \\
&= \sum_{(S,F)\in \mathrm{Correct}_{G_n}(A)} \frac{P_{G_n}{}'((S,F))}{p^{|ISO_{G_n}\cap F|}(1-p)^{|ISO_{G_n}\cap(U_n-F)|}} \\
&\geq \frac{\sum_{(S,F)\in \mathrm{Correct}_{G_n}(A)} P_{G_n}{}'((S,F))}{[\max(p, 1-p)]^{|ISO_{G_n}|}}
\end{aligned}$$

and thus

$$\begin{aligned}
&P_{G_n}{}'(\mathrm{NotCorrect}_{G_n}(A)) \\
&\geq \left(\frac{1}{[\max(p,1-p)]^{|ISO_{G_n}|}} - 1\right) \sum_{(S,F)\in \mathrm{Correct}_{G_n}(A)} P_{G_n}{}'((S,F))
\end{aligned}$$

So, by assumption $\forall P_{G_n}' \in \mathcal{P}_{G_n}$ and $\forall n \geq n_0$

$$P_{G_n}'(\text{NotCorrect}_{G_n}(A)) \geq \frac{1 - [\max(p, 1-p)]^{|ISO_{G_n}|}}{[\max(p, 1-p)]^{|ISO_{G_n}|}} \cdot c$$
$$\rightarrow \infty$$

This is clearly a contradiction, implying that for any deterministic diagnosis algorithm A and $\forall P_{G_n}' \in \mathcal{P}_{G_n}$,

$$P_{G_n}'(\text{Correct}_{G_n}(A)) \rightarrow 0.$$

Thus, for any algorithm A

$$\text{PCD}_{G_n}(A) \rightarrow 0$$

as well. Now, consider any random diagnosis algorithm A. Then, $\forall P_{G_n}' \in \mathcal{P}_{G_n}$

$$\text{PCD}_{G_n}(A) \leq \sum_{(S,F) \in \Omega_{G_n}} P_{G_n}'((S, F)) \cdot p_{A,S}(F)$$

Consider the deterministic algorithm A' that for any syndrome S chooses fault set F such that $\forall F' \subseteq U_n$

$$P_{G_n}'((S, F)) \geq P_{G_n}'((S, F')).$$

Then, if $\mathcal{S}$ represents the set of all syndromes in G_n

$$\begin{aligned}
\text{PCD}_{G_n}(A) &\leq \sum_{(S,F) \in \Omega_{G_n}} P_{G_n}'((S, Faulty_{A'}(S))) \cdot p_{A,S}(F) \\
&= \sum_{S \in \mathcal{S}} \sum_{F \subseteq U_n} P_{G_n}'((S, Faulty_{A'}(S))) \cdot p_{S,S}(F) \\
&= \sum_{S \in \mathcal{S}} P_{G_n}'((S, Faulty_{A'}(S))) \sum_{F \subseteq U_n} p_{A,S}(F) \\
&= P_{G_n}'(\text{Correct}_{G_n}(A')) \\
&\rightarrow 0
\end{aligned}$$

∎

This proof, in fact, yields a stronger result than is stated in Theorem 3 in the following sense. The probability of correct diagnosis is defined as the minimum probability over all possible behaviors of the faulty processors. The result given in Theorem 3 is shown to be true for all refinements, meaning that no matter how the faulty processors act, no algorithm will be capable of achieving correct diagnosis with high probability.

8 Diagnosis in Regular Systems

The study of regular systems is important for several reasons. First, in many application areas such as VLSI circuitry, regular designs are easily and efficiently implementable. Furthermore, the majority of existing multiprocessor systems possess a regular structure. Finally, the maximum number of tests conducted by any processor is one measure of the overhead required to achieve fault tolerance. For a fixed total number of tests, regular systems require the minimum overhead under this measure.

The diagnosis algorithm given in [5] was shown to achieve correct diagnosis with probability approaching one in a class of regular systems known as $D_{1,k}$ systems that conduct $cn \log n$ tests, for $c > 1/(\log 1/p)$. Furthermore, it was proven in [10] under a more general probability model that all diagnosis algorithms must have probability approaching zero of correct diagnosis in regular systems where the number of tests grows more slowly than $n \log n$. This more general probability model contains the model utilized in this paper as a special case and hence this result holds for this model as well. Thus, for the important class of regular systems the algorithm given in [5] is optimal to within a constant factor. This also demonstrates that the irregular structure of the tester digraphs studied in this paper are a crucial factor in making them amenable to diagnosis.

9 Conclusion

A uniformly probabilistic fault model for multiprocessor systems in which processors are faulty with probability p has been studied. It has been shown that correct diagnosis can be achieved with probability approaching one in a class of systems that conducts slightly more than a linear number of tests using a simple and efficient diagnosis algorithm. It has also been shown that this result is very close to the best possible, i.e. that in systems conducting a number of tests that grows more slowly than n all diagnosis algorithms, whether they be deterministic or random must have a probability approaching zero of correct diagnosis.

References

[1] Hakimi, S. L., and A. T. Amin, "Characterization of the Connection Assignment of Diagnosable Systems," *IEEE Transactions on Computers*, vol. C-23, pp. 86–88, January 1974.

[2] Maheshwari, S. N., and S. L. Hakimi, "On Models for Diagnosable Systems and Probabilistic Fault Diagnosis," *IEEE Transactions on Computers,* vol. C-25, pp. 228–236, March 1976.

[3] Sullivan, G. F., "System-Level Fault Diagnosability in Probabilistic and Weighted Models," *Digest of the 17th International Symposium on Fault-Tolerant Computing,* IEEE Computer Society Press, pp. 190–195, 1987.

[4] Dahbura, A. T., "An Efficient Algorithm for Identifying the Most Likely Fault Set in a Probabilistically Diagnosable System," *IEEE Transactions on Computers,* vol. C-35, pp. 354-356, April 1986.

[5] Scheinerman, E., "Almost Sure Fault Tolerance in Random Graphs," *SIAM Journal on Computing,* vol. 16, pp. 1124-1134, December 1987.

[6] Blough, D. M., G. F. Sullivan, and G. M. Masson, "Almost Certain Diagnosis for Intermittently Faulty Systems," *Digest of the 18th International Symposium on Fault-Tolerant Computing,* IEEE Computer Society Press, pp. 260–265, 1988.

[7] Preparata, F., G. Metze, and R. Chien, "On the Connection Assignment Problem of Diagnosable Systems," *IEEE Transactions on Electronic Computers,* vol. EC–16, pp. 848-854, December 1967.

[8] Angluin, D., and L. Valiant, "Fast Probabilistic Algorithms for Hamiltonian Circuits and Matchings," *Journal of Computer and System Sciences,* vol. 18, pp. 155–193, April 1979.

[9] Chernoff, H., "A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the Sum of Observations," *Annals of Mathematical Statistics,* vol. 23, pp. 493–507, 1952.

[10] Blough, D. M., *Fault Detection and Diagnosis in Multiprocessor Systems,* Ph. D. Dissertation, The Johns Hopkins University, 1988.